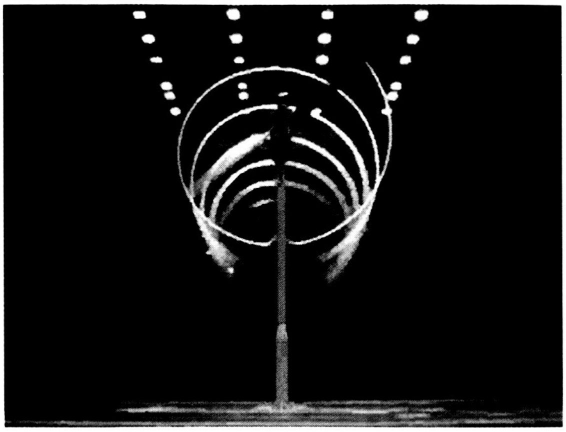

Fig. 2. The research wind turbine inside the 80- by 120-Foot Wind Tunnel expelling smoke for wake visualization; airspeed = 7 m/sec, blade pitch = 3 degrees, yaw angle = 0.0 degrees.

the research wind turbine in the NASA Ames Full-Scale Aerodynamic Complex (80- by 120-Foot Wind Tunnel) (Fig. 2).

The research wind turbine was 10 meters (m) in diameter, turned at 72 revolutions per minute (rpm), and could generate 20 kilowatts (kW) of power. The turbine was highly instrumented with over 150 pressure transducers, strain gages, and motion sensors to identify its operational state. Easily reconfigured with both blade-pitch and nacelle azimuth control, the research wind turbine was tested in many different operational variations. Nominal testing was at airspeeds between 5 and 25 meters per second (m/sec) with a few test points recorded at 40 m/sec.

Point of Contact R. Kufeld
(650) 604-5664
rkufeld@mail.arc.nasa.gov

## Remote Access and Analysis of Aeronautics Data

David Korsmeyer, Joan Walton

The DARWIN system has been developed in conjunction with new computational and experimental test technologies to provide remote access to the integrated knowledge generated from these independent systems. Version 2.5 of the DARWIN system for remote access to and analysis of aeronautics data was completed and deployed at the Ames wind tunnels in December 1999. The new and improved version was used by the Wind Tunnel Operations Division to support one of the first tests in the newly modernized 11- by 11-Foot Transonic Wind Tunnel in early January 2000.

The deployment of DARWIN version 2.5 at the Ames wind tunnels concluded a year of work by the DARWIN development team on implementing enhancements to the DARWIN system. The major new feature of this release is a completely redesigned database schema. Building on 3 years of experience with the original database and on feedback from collaborators at NASA Langley, Eglin Air Force Base, and Arnold Engineering Development Center, the team developed a schema that can store values for a wide range of variables and thus flexibly capture information from disparate wind-tunnel test entries. All components of the DARWIN system, including the client, administrator, and database loader, were completely reworked to take advantage of the new database schema.

In addition to the new database, other features of this release include an improved client user interface and ability to review data from additional tunnel instrumentation suites, such as Video Model Deformation (VMD), Temperature Sensitive Paint (TSP), and Mini Tufts (see fig. 1). It also includes the capability to manage and cross-plot associated computation fluid dynamics (CFD) results with the experimental data. The new DARWIN administrator allows wind-tunnel support personnel to manage the system without requiring assistance from the development team. Finally, the automated database loading system now takes a standard file format (netCDF) as input so that accepting information produced by other data acquisition systems is easier.

Completion of the 11-foot LB435 Calibration Test in the 11- by 11-Foot Transonic Wind Tunnel marked the first support of a live wind-tunnel test entry using DARWIN version 2.5. All aspects of the system performed as advertised, including the client application, the

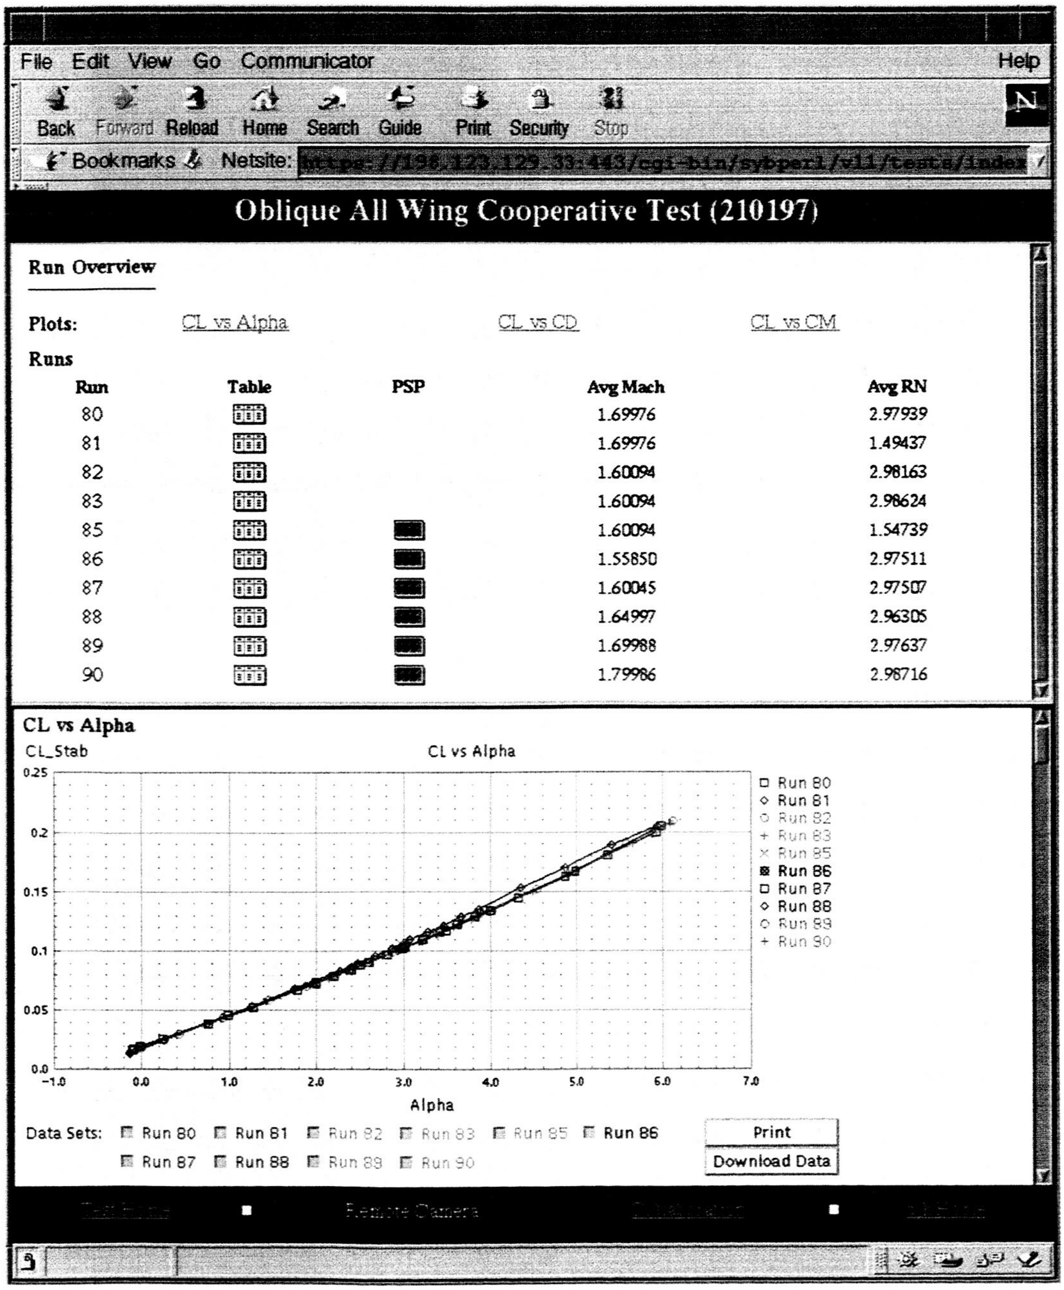


| Run | Table | PSP | Avg Mach | Avg RN |
|---|---|---|---|---|
| 80 | | | 1.69976 | 2.97939 |
| 81 | | | 1.69976 | 1.49437 |
| 82 | | | 1.60094 | 2.98163 |
| 83 | | | 1.60094 | 2.98624 |
| 85 | | | 1.60094 | 1.54739 |
| 86 | | | 1.55850 | 2.97511 |
| 87 | | | 1.60045 | 2.97507 |
| 88 | | | 1.64997 | 2.96305 |
| 89 | | | 1.69988 | 2.97637 |
| 90 | | | 1.79986 | 2.98716 |

*Fig. 1. The Darwin system screen showing the oblique all wing cooperative test (210197).*

automated database loader, and the administrator application. The Wind Tunnel Operations staff made a few feature requests, which were addressed and implemented within a week or two of receipt. During the LB435 test, values for 119 variables were collected at 19,684 points in time, resulting in approximately 2.2 million recorded values. The DARWIN database now holds close to 20 million values across 18 wind-tunnel tests.

Deployment of DARWIN version 2.5 at Ames allows wind-tunnel customers, including NASA and university researchers, the U.S. military, and commercial aircraft manufacturers, to access wind-tunnel test data from remote locations. Test data are available in near real time while tests are in progress and can be compared against archival tests already in the database. The DARWIN system lowers the cost of wind-tunnel tests by reducing travel requirements and providing rapid data analysis capabilities. Advances made in the version 2.5 release facilitate DARWIN support of the Unitary Plan Wind Tunnels and the National Full-Scale Aerodynamics Complex, and also pave the way for adoption of the DARWIN system at NASA Langley Research Center.

Point of Contact: David Korsmeyer
(650) 604-3144
dkorsmeyer@mail.arc.nasa.gov

## Remote Large Data Set Visualization

David Ellsworth

Simulations run on large parallel systems produce data sets that contain hundreds of megabytes to terabytes of data. The researchers producing these data sets prefer to visualize them using their personal workstations. High-end PC workstations currently have the computation and graphics power to perform these visualizations. However, these workstations do not have sufficient memory to completely load large data sets. This research, which enables the use of personal workstations for visualizing data sets that are too large to be stored on personal workstations, will increase the productivity of researchers who are pushing the limits of simulations by producing very large data sets.

Because personal workstations have limited memory, out-of-core visualization techniques must be used. These techniques calculate the visualization with only a fraction of the data set resident in memory. In addition, many data sets are so large that they can only fit on central file servers. Since most file servers do not have significant extra central processing unit (CPU) and memory capacity, remote out-of-core visualization is required.

Our earlier research developed an out-of-core visualization technique called application-controlled demand paging. This technique loads the data required by the visualization algorithm from disk into main memory as necessary. The technique works well because most visualization algorithms only use a small fraction of the entire data set. The overall speed typically increases as the user interacts with the data set, since previously loaded data are retained if possible. This means that the overall speed will soon approach the speed that would be seen if the entire data set were loaded into memory.

However, the original implementation of out-of-core visualization using demand paging did not try to perform computation, disk access, or network access at the same time. When the implementation detected that data must be